

I hereby certify that this correspondence  
is being deposited with the United States  
Postal Service with sufficient postage as  
first class mail in an envelope addressed to:

Commissioner for Patents  
P.O. Box 1450  
Alexandria, VA 22313-1450

*Richard J. Coldren*  
Attorney for Applicant

Date: 11 March 2004

Application No. : 10/689,355  
Applicant : Mark Beaumont  
Filed : 20 October 2003  
Title : METHOD FOR USING FILTERING TO LOAD BALANCE A LOOP  
OF PARALLEL PROCESSING ELEMENTS  
  
TC/A.U. : Not yet assigned  
Examiner : Not yet assigned  
  
Atty. Docket No. : DB001061-000  
Customer No. : 24122

Commissioner for Patents  
P.O. Box 1450  
Alexandria, VA 22313-1450

#### TRANSMITTAL LETTER

Dear Sir:

Enclosed for filing in the above-captioned application, please find the following:

- Submission of Foreign Priority Document UK Patent Application No. 0309202.0 as required Under 35 U.S.C. §119 and 37 C.F.R. §1.55.

Also enclosed is a return postcard. Please date stamp the postcard and return it to us to acknowledge receipt of this patent application.

The Commissioner is hereby authorized to charge any underpayment or credit any overpayment to our Deposit Account No. 20-0888. A duplicate copy of this Transmittal letter is enclosed.

Respectfully submitted,

*Richard J. Coldren*

Richard J. Coldren  
Reg. No. 44,084  
THORP REED & ARMSTRONG, LLP  
One Oxford Centre, 14<sup>th</sup> Floor  
301 Grant Street  
Pittsburgh, PA 15219-1425  
(412) 394-2442

Dated: 11 March 2004

Attorney for Applicant





INVESTOR IN PEOPLE

The Patent Office  
Concept House  
Cardiff Road  
Newport  
South Wales  
NP10 8QQ

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation and Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein together with the Statement of inventorship and of right to grant of a Patent (Form 7/77), which was subsequently filed.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.


Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.



Signed

Dated 27 February 2004





1/77  
24APR03 102074-23 000389  
P01/77000.00-0309202.0

# Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

The Patent Office

Cardiff Road  
Newport  
South Wales  
NP10 8QQ

1. Your reference

P17575GB-JCC/lp

2. Patent application number

(The Patent Office will fill in this part)

0309202.0

3. Full name, address and postcode of the or of each applicant (underline all surnames)

Micron Europe Limited,  
Micron House,  
Wellington Business Park,  
Dukes Ride, Crowthorne,  
Berkshire RG45 6LS

Patents ADP number (if you know it)

8291817002

If the applicant is a corporate body, give the country/state of its incorporation

4. Title of the invention

METHOD FOR USING FILTERING TO LOAD BALANCE A LOOP OF PARALLEL PROCESSING ELEMENTS

5. Name of your agent (if you have one)

Forrester Ketley &amp; Co.

"Address for service" in the United Kingdom to which all correspondence should be sent (including the postcode)

Forrester House  
52 Bounds Green Road  
London  
N11 2EY

Patents ADP number (if you know it)

133001 ✓

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number

Country

Priority application number  
(if you know it)

Date of filing  
(day / month / year)

GB  
GB

0221562.2  
0221563.0

17/09/2002  
17/09/2002

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application

Date of filing  
(day / month / year)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:

YES

- a) any applicant named in part 3 is not an inventor, or
  - b) there is an inventor who is not named as an applicant, or
  - c) any named applicant is a corporate body.
- See note (d))

# Patents Form 1/77

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form

Description	14
Claim(s)	7
Abstract	1
Drawing(s)	5

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (Patents Form 7/77)

Request for preliminary examination and search (Patents Form 9/77) 1

Request for substantive examination (Patents Form 10/77)

Any other documents (please specify)

11. I/We request the grant of a patent on the basis of this application.

Signature *Forrester Ketley & Co* Date 22 April 2003

Forrester Ketley & Co.


12. Name and daytime telephone number of person to contact in the United Kingdom (020) 8889 6622 J.C. Carstairs

## Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

## Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 08459 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.



7/77 04  
06MAY03 8130950-2 000389  
P07/7700 0-00-0309202.0

**Statement of inventorship and of  
right to grant of a patent**

The Patent Office

Cardiff Road  
Newport  
South Wales  
NP10 8QQ

1. Your reference	P17575GB-JCC/lp
2. Patent application number (if you know it)	0309202.0
3. Full name of the or of each applicant	Micro Europe Limited
4. Title of the invention	METHOD FOR USING FILTERING TO LOAD BALANCE A LOOP OF PARALLEL PROCESSING ELEMENTS
5. State how the applicant(s) derived the right from the inventor(s) to be granted a patent	As employers of the inventor
6. How many, if any, additional Patents Forms 7/77 are attached to this form? (see note (c))	
7.	I/We believe that the person(s) named over the page (and on any extra copies of this form) is/are the inventor(s) of the invention which the above patent application relates to.  <div style="display: flex; justify-content: space-between;"> <div> <p>Signature <i>Forrester Ketley &amp; Co.</i> Forrester Ketley &amp; Co.</p> </div> <div> <p>Date 2 May 2003</p> </div> </div>
8. Name and daytime telephone number of person to contact in the United Kingdom	(020) 8889 6622 CARSTAIRS, J.C.

**Notes**

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 08459 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there are more than three inventors, please write the names and addresses of the other inventors on the back of another Patents Form 7/77 and attach it to this form.
- When an application does not declare any priority, or declares priority from an earlier UK application, you must provide enough copies of this form so that the Patent Office can send one to each inventor who is not an applicant.
- Once you have filled in the form you must remember to sign and date it.

Enter the full names, addresses and postcodes of the inventors in the boxes and underline the surnames

Mark BEAUMONT,  
30 Lorne Street,  
Reading,  
RG1 7YN  
Berkshire,  
United Kingdom

8623191001

Patents ADP number (if you know it):

Patents ADP number (if you know it):

Patents ADP number (if you know it):

Reminder

Have you signed the form?



SECRET

## METHOD FOR USING FILTERING TO LOAD BALANCE A LOOP OF PARALLEL PROCESSING ELEMENTS

### BACKGROUND OF THE INVENTION

[0002] The present invention relates generally to parallel processing and more particularly to balancing the work loads of the processing elements within a parallel processing system.

[0003] Conventional central processing units ("CPU's"), such as those found in most personal computers, execute a single program (or instruction stream) and operate on a single stream of data. For example, the CPU fetches its program and data from a random access memory ("RAM"), manipulates the data in accordance with the program instructions, and writes the results back sequentially. There is a single stream of instructions and a single stream of data (note: a single operation may operate on more than one data item, as in  $X = Y + Z$ , however, only a single stream of results is produced). Although the CPU may determine the sequence of instructions executed in the program itself, only one operation can be completed at a time. Because conventional CPUs execute a single program (or instruction stream) and operate on a single stream of data, conventional CPUs may be referred to as a single-instruction, single data CPU or an SISD CPU.

[0004] The speed of conventional CPUs has dramatically increased in recent years. Additionally, the use of cache memories enables conventional CPUs faster access to the desired instruction and data streams. However because conventional CPUs can only complete

one operation at a time, conventional CPUs are not suitable for extremely demanding applications having large data sets (such as moving image processing, high quality speech recognition, and analytical modeling applications, among others).

[0005] Improved performance over conventional SISD CPUs may be achieved by building systems which exhibit parallel processing capability. Typically, parallel processing systems use multiple processing units or processing elements to simultaneously perform one or more tasks on one or more data streams. For example in one class of parallel processing system, the results of an operation from a first CPU are passed to a second CPU for additional processing, and from the second CPU to another CPU, and so on. Such a system, commonly known as a "pipeline", is referred to as a multiple-instruction, single-data or MISD system because each CPU receives a different instruction stream while operating on a single data stream. Improved performance may also be obtained by using a system which contains many autonomous processors, each running its own program (even if the program running on the processors is the same code) and producing multiple data streams. Systems in this class are referred to as a multiple-instruction, multiple-data or MIMD system.

[0006] Additionally, improved performance may be obtained using a system which has multiple identical processing units each performing the same operations at once on different data streams. The processing units may be under the control of a single sequencer running a single program. Systems in this class are referred to as a single-instruction, multiple data or SIMD system. When the number of processing units in this type of system is very large (e.g., hundreds or thousands), the system may be referred to as a massively parallel SIMD system.

[0007] Nearly all computer systems now exhibit some aspect of one or more of these types of parallelism. For example, MMX extensions are SIMD; multiple processors (graphics processors, etc) are MIMD; pipelining (especially in graphics accelerators) is MISD. Furthermore, techniques such as out of order execution and multiple execution units have been used to introduce parallelism within conventional CPUs as well.

[0008] Parallel processing is also used in active memory applications. An active memory refers to a memory device having a processing resource distributed throughout the memory structure. The processing resource is most often partitioned into many similar processing elements (PEs) and is typically a highly parallel computer system. By distributing the processing resource throughout the memory system, an active memory is able to exploit the very high data bandwidths available inside a memory system. Another advantage of active memory is that data can be processed "on-chip" without the need to transmit the data across a system bus to the CPU or other system resource. Thus, the work load of the CPU may be reduced to operating system tasks, such as scheduling processes and allocating system resources.

[0009] A typical active memory includes a number of interconnected PEs which are capable of simultaneously executing instructions sent from a central sequencer or control unit. The PEs may be connected in a variety of different arrangements depending on the design requirements for the active memory. For example, PEs may be arranged in hypercubes, butterfly networks, one-dimensional strings/loops, and two-dimensional meshes, among others.

[0010] In typical active memories, load imbalances often occur such that some PEs are idle (i.e., without assigned tasks) while other PEs have multiple tasks assigned. To maximize the effectiveness of the active memory, it is desirable to balance the work load across all of the PEs. For example in an active memory having a multitude of identical PEs, it is desirable that each PE be assigned the same number of instructions by the central sequencer, thus maximizing the resources of the active memory. Additionally in an active memory having non-identical PEs, it may be desirable to assign more tasks to the PEs with greater processing capabilities. By balancing the load, the amount of time that one or more PEs is idle while waiting for one or more other PEs to complete their assigned tasks is minimized.

[0011] Thus, there exists a need for a method for balancing the load of a parallel processing system such that the resources of the parallel processing system are maximized. More specifically, there exists a need for a method for balancing the load of an active memory such that the resources of the active memory are maximized.

#### SUMMARY OF THE INVENTION

[0012] One aspect of the present invention relates to a method for balancing the load of a parallel processing system having a plurality of parallel processing elements arranged in a loop, wherein each processing element has a local number of tasks ( $v_r$ ) associated therewith, wherein  $r$  represents the number for a selected processing element  $PE_r$ , and wherein each of the processing elements are operable to communicate with a clockwise adjacent processing element and with an anti-clockwise adjacent processing element. The method comprises determining within each processing element ( $PE_r$ ) a total number of tasks ( $V$ ) present within the loop, calculating a local mean number of tasks ( $M_r$ ) within each of the plurality of processing elements ( $PE_r$ ), calculating a local deviation ( $D_r$ ) within each of the plurality of processing elements ( $PE_r$ ). The method also comprises determining a sum weighted deviation within each of the processing elements ( $PE_r$ ) for one-half the loop in an anti-clockwise direction ( $A$ ), the one-half of the loop being relative to each of the selected processing elements ( $PE_r$ ), determining a sum weighted deviation within each of the processing elements ( $PE_r$ ) for one-half the loop in a clockwise direction ( $C$ ), the one-half of the loop being relative to each of the selected processing elements ( $PE_r$ ); determining a clockwise transfer parameter

( $T_c$ ) and an anti-clockwise transfer parameter ( $T_a$ ) within each of the processing elements ( $PE_r$ ), and redistributing tasks among the plurality of processing elements in response to the clockwise transfer parameters ( $T_c$ ) and the anti-clockwise parameters ( $T_a$ ) within each of the plurality of processing elements ( $PE_r$ ).

[0013] Another aspect of the present invention relates to a method for assigning tasks among a plurality of processing elements within a parallel processing system, the processing elements being connected in a loop and having a local number of tasks ( $v_r$ ) associated therewith. The method comprises determining the total number of tasks on the loop, computing a local mean value for each of the processing elements, assigning a weight to each of said plurality of processing elements, and computing a local weighted deviation for each of the processing elements, the local deviation representative of the difference between the local number of tasks for a processing element and the local mean value for the processing element. The method also includes summing the weighted deviation of the processing elements located within one-half of the loop in an anti-clockwise direction, summing the weighted deviation of the processing elements located within one-half of the loop in a clockwise direction, computing a number of tasks to transfer in a clockwise direction, computing a number of tasks to transfer in an anti-clockwise direction, and redistributing tasks relative to the number of tasks to transfer in a clockwise direction and the number of tasks to transfer in an anti-clockwise direction.

[0014] The present invention enables tasks to be distributed along a group of serially connected PEs so that each PE typically has  $X$  number of tasks or  $(X+1)$  number of tasks to perform in the next phase. The present invention may be performed using the hardware and software (i.e., the local processing capability) of each PE within the array. Those advantages and benefits, and others, will become apparent from description of the invention below.

#### BRIEF DESCRIPTION OF THE DRAWINGS

[0015] To enable the present invention to be easily understood and readily practiced, the present invention will now be described for purposes of illustration and not limitation, in connection with the following figures wherein:

[0016] FIG. 1 is a block diagram illustrating an active memory according to an embodiment of the present invention.

[0017] FIG. 2 is a block diagram of a processing element for the active memory as illustrated in FIG. 1 according to an embodiment of the present invention.

[0018] FIG. 3 illustrates an array of the processing elements as illustrated in FIG. 2 arranged in a loop according to an embodiment of the present invention.

[0019] FIG. 4 illustrates an operational process for balancing the load within a loop of processing elements according to various embodiments of the present invention.

[0020] FIG. 5 illustrates the determination of the sum weighted deviation in the anti-clockwise half of loop (A) and the determination of the sum weighted deviation in the clockwise half of loop (C) for a local PE according to an embodiment of the present invention.

#### DETAILED DESCRIPTION OF THE INVENTION

[0021] As discussed above, parallel processing systems may be placed within one or more classifications (e.g., MISD, MIMD, SIMD, etc.). For simplicity, the present invention is discussed in the context of a SIMD parallel processing system. More specifically, the present invention is discussed in the context of a SIMD active memory. It should be noted that such discussion is for clarity only and is not intended to limit the scope of the present invention in any way. The present invention may be used for other types and classifications of parallel processing systems.

[0022] FIG. 1 is a block diagram illustrating an active memory 10 according to an embodiment of the present invention. It should be noted that the active memory 10 is only one example of a device on which the methods of the present invention may be practiced and those of ordinary skill in the art will recognize that the block diagram of FIG. 1 is an overview of an active memory device 10 with a number of components known in the art being omitted for purposes of clarity.

[0023] Active memory 10 is intended to be one component in a computer system. Processing within active memory 10 is initiated when the active memory 10 receives commands from a host processor (not shown), such as the computer system's CPU. A complete processing operation (i.e., data movement and processing) in the active memory 10 may consist of a sequence of many commands from the host to the active memory device 10.

[0024] Active memory 10 is comprised of a host memory interface ("HMI") 12, a bus interface 14, a clock generator 16, a task dispatch unit ("TDU") 18, a DRAM control unit ("DCU") 20, a DRAM module 22, a programmable SRAM 24, an array control sequencer 26, and a processing element array 28, among others.

[0025] The HMI 12 provides an input/output channel between the host (such as a CPU, not shown) and the DRAM module 22. In the current embodiment, the HMI 12 receives command (cmd), address (addr), and data signals (among others) from and sends data and ready (rdy) signals (among others) to the host. The HMI 12 approximates the operation of a standard non-active memory so that the host, without modifications, is compatible with the active memory 10.

[0026] The HMI 12 may be similar in its operation to the interface of a synchronous DRAM as is known in the art. Accordingly, the host must first activate a page of data to access data within a DRAM module 22. In the current embodiment, each page may contain 1024 bytes of data and there may be 16,384 pages in all. Once a page has been activated, it can be written and read through the HMI 12. The data in the DRAM module 22 may be updated when the page is deactivated. The HMI 12 also sends control signals (among others) to the DCU 20 and to the processing element array 28 via the task dispatch unit 18.

[0027] The HMI 12 may operate at a frequency different than that of the frequency of the master clock. For example, a 2x internal clock signal from clock generator 16 may be used. Unlike a traditional DRAM, the access time for the HMI 12 uses a variable number of cycles to complete an internal operation, such as an activate or deactivate. Thus the ready signal (rdy) is provided to allow the host to detect when a specific command has been completed.

[0028] The bus interface 14 provides an input/output channel between the host and the TDU 18. For example, the bus interface 14 receives column select (cs), write command (w), read command (r), address (addr), and data signals (among others) from and places interrupt (intr), flag, and data signals (among others) onto the system bus (not shown). The bus interface 14 also receives signals from and sends signals to TDU 18.

[0029] The clock generator 16 is operable to receive an external master clock signal (x1) and operable to provide the master clock signal (x1) and one or more internal clock signals (x2, x4, x8) to the components of the active memory. It should be apparent to one skilled in the art that other internal clock signals may be produced by the clock generator 16.

[0030] The TDU 18 communicates with the bus interface 14, the HMI 12, the programmable SRAM 24, the array control sequencer 26, and the DCU 20. In the current embodiment, the TDU 18 functions as an interface to allow the host to issue a sequence of commands to the array control sequencer 26 and the DCU 20. Task commands from the host may be buffered in the TDU's FIFO buffers to allow a burst command to be issued. Commands may contain information on how the tasks in the array control sequencer 26 and the DCU 20 should be synchronized with one another, among others.

[0031] The DCU 20 arbitrates between the TDU 18 and the HMI 12 and sends commands to the DRAM modules 22 and the processing element array 28. The DCU 20 also schedules refreshes within the DRAM modules 22. In one embodiment, the DRAM modules 22 of the active memory 10 may be comprised of sixteen 64k x128 eDRAM (or embedded DRAM) cores. Each eDRAM core may be connected to an array of sixteen PEs, thus providing 256 (16 x 16) PEs in all.

[0032] The programmable SRAM 24 functions as a program memory by storing commands issued by the TDU 18. For example, the TDU 18 may transmit a "write program memory

address" command which sets up a start address for a write operation and a "write program memory data" command which writes a memory location and increments the program memory write address, among others. The programmable SRAM 24, in the current embodiment, has both an address register and a data output register.

[0033] The array control sequencer 26 may be comprised of a simple 16 bit minimal instruction set computer (16-MISC). The array control sequencer 26 communicates with the TDU 18, the programmable SRAM 24, and the DCU 20, and is operable to generate register file addresses for the processing element array 28 and operable to sequence the array commands, among others.

[0034] The processing element array 28 is comprised of a multitude of processing elements ("PEs") 30 (see FIG. 2) connected in a variety of different arrangements depending on the design requirements for the processing system. For example, processing units may be arranged in hypercubes, butterfly networks, one-dimensional strings/loops, and two-dimensional meshes, among others. In the current embodiment, the processing elements 30 are arranged in a loop (for example, see FIG. 3). The processing element array 28 communicates with the DRAM module 22 and executes commands received from the programmable SRAM 24, the array control sequencer 26, the DCU 20, and the HMI 12. Each PE in the processing element array 28 includes dedicated H-registers for communication with the HMI 12. Control of the H-registers is shared by the HMI 12 and the DCU 20.

[0035] Referring now to FIG. 2, a block diagram of a PE 30 according to one embodiment of the present invention is illustrated. PE 30 includes an arithmetic logic unit ("ALU") 32, Q-registers 34, M-registers 36, a shift control and condition register 38 (also called "condition logic" 38), a result register pipeline 40, and register file 42. The PE 30 may also contain other components such as multiplexers 46 and logic gates (not shown), among others.

[0036] In the current embodiment, the Q-registers 34 are operable to merge data into a floating point format and the M-Registers 36 are operable to de-merge data from a floating point format into a single magnitude plus an exponent format. The ALU 32 is a multiplier-adder operable (among others) to receive information from the Q-registers 34 and M-registers 36, execute tasks assigned by the TDU 18 (see FIG. 1), and transmit results to the shift control and condition logic 38 and to the result register pipeline 40. The result register pipeline 40 is operable to communicate with the register file 42, which holds data for transfer into or out of the DRAM modules 22 via a DRAM interface 44. Data is transferred between the PE and the DRAM module 22 via a pair a registers, one register being responsive to the DCU 20 and the other register being responsive to the PE 30. The DRAM interface receives command information from the DCU 20. The DRAM interface 44 also permits the PE 30 to communicate with the host through the host memory access port 46.

[0037] In the current embodiment, the H-registers 42 are comprised of synchronous SRAM and each processing element within the processing element array 28 contains eight H-registers 42 so that two pages can be stored from different DRAM locations, thus allowing the interleaving of short i/o bursts to be more efficient. Result register pipeline 40 is also connected to one or more neighborhood connection registers ("X-register") (not shown). The X-register links one PE 30 to its neighboring PE's 30 in the processing element array 28.

[0038] The reader desiring more information about the hardware shown in FIGs. 1 and 2 is directed to UK Patent application (serial no. not yet assigned) entitled "Control of Processing Elements in Parallel Processors" filed 17 September 2002, (Micron no. 02-1604) which is hereby incorporated by reference. Details about the PEs may also be found in UK Patent Application No. 021562.2 entitled "Host Memory Interface for a Parallel Processor" filed 17 September 2002, (Micron no. 02-0703) which is hereby incorporated by reference.

[0039] FIG. 3 is a simplified diagram showing the interconnections of an array of the processing elements 30 (as illustrated in FIG. 2) arranged in a loop 50 according to an embodiment of the present invention. In the current embodiment, loop 50 is comprised of eight (8) PEs 30 (i.e.,  $PE_0, PE_1, \dots, PE_7$ ) which are interconnected via their associated X-register links. It should be noted that the number of PEs 30 included in loop 50 may be altered while remaining within the scope of the present invention. As illustrated in FIG. 3, each PE is operable to communicate with its clockwise and anti-clockwise neighbor. For example,  $PE_1$  is operable to communicate with its clockwise neighbor,  $PE_2$ , and with its anti-clockwise neighbor,  $PE_0$ . In the current embodiment, every PE 30 on the loop 50 receives instructions from a single TDU 18 as discussed in conjunction with FIG. 1. Furthermore, each PE has a local number of tasks ( $v_r$ ) associated therewith. For example,  $PE_0$  has three (3) tasks associated therewith (i.e.,  $v_0 = 3$ ),  $PE_1$  has six (6) tasks associated therewith (i.e.,  $v_1 = 6$ ),  $PE_2$  has two (2) tasks associated therewith (i.e.,  $v_2 = 2$ ), etc.

[0040] FIG. 4 illustrates an operational process 60 for balancing the work loads between the PEs 30 on loop 50 according to an embodiment of the present invention. Operational process 60 begins by determining the total number of tasks ( $V$ ) present on the loop in operation 61. As discussed above in conjunction with FIG. 3, each  $PE_r$  (where  $r = 0$  to 7, e.g.,  $PE_0, PE_1, \dots, PE_7$ ) in the loop has a local number of tasks ( $v_r$ ) associated therewith. In the current embodiment, each  $PE_r$  passes its own value  $v_r$  onto its clockwise neighbor and simultaneously receives a value  $v_{r+1}$  from its anti-clockwise neighbor. Each  $PE_r$  keeps a running partial sum (i.e., adds each value  $v_{r+1}$  received to its own value  $v_r$ ). This process continues until each value  $v_r$  has moved clockwise around the loop and visited each  $PE_r$ , in this case seven transfers are needed. At the end of the rotation process, the sum represents the total number



of tasks ( $V$ ) on the loop. As illustrated in FIG. 3, loop 50 has forty-three (43) total tasks associated therewith.

[0041] The sum ( $V$ ) can be expressed by the equation  $V = \sum_{i=0}^{N-1} v_i$ , where  $N$  represents the number of PEs 30 in the loop 50 (here  $N = 8$ ), and  $v_i$  represents the local number of tasks associated with an  $i^{\text{th}}$  processing element in the loop. For example, for  $i = 3$ , the number of tasks associated with  $PE_3$  (i.e.,  $v_3$ ) is added to the sum  $V$ . It should be noted that after a rotation is completed, each  $PE_r$  will have calculated the same value for ( $V$ ). It should also be noted in the current discussion, "local" refers to the values or functions associated with a single PE within the loop, whereas "global" refers to the values or function associated with the entire loop of PEs.

[0042] After the total number of tasks ( $V$ ) present on the loop is determined in operation 61, the local mean number ( $M_r$ ) of tasks for each  $PE_r$  is computed in operation 62. In the current embodiment, operation 62 employs a rounding function to ensure that no tasks are lost or

"gained" during the rounding process (i.e., to ensure that  $V = \sum_{i=0}^{N-1} M_i$ ).

[0043] For example assume that 13 tasks (i.e.,  $V = 13$ ) are to be shared by the eight PEs (i.e.,  $PE_0$  through  $PE_7$ ). Without the rounding function, the local mean for each PE would be  $PE_r = 1.625$  before rounding (i.e.,  $13 \div 8 = 1.625$ ). If the fraction thirteen-eighths is set to round down for each PE (i.e.,  $13 \div 8 = 1$ ), then the sum of the means for all of the individual PEs (i.e.,  $PE_0$  through  $PE_7$ ) is equal to eight (8) and five ( $13 - 8 = 5$ ) tasks are lost. In contrast, if the fraction thirteen-eighths is set to round up for each PE (i.e.,  $13 \div 8 = 2$ ), then the sum of the means for all of the individual PEs (i.e.,  $PE_0$  through  $PE_7$ ) is equal to sixteen (16) and three ( $16 - 13 = 3$ ) extra tasks are gained. The rounding function is discussed in more detail in U.S. Patent Application Serial No. \_\_\_\_\_ entitled "Method for Rounding Values for a Plurality of Parallel Processing Elements" filed \_\_\_\_\_ (DB001064-000, Micron no. 02-1269) and incorporated in its entirety by reference herein.

[0044] The rounding function  $M_r = \text{Trunc}((V + E_r) / N)$  prevents tasks from being lost or gained (where  $M_r$  represents the local mean for  $PE_r$ ,  $2N$  represents the total number of PEs 30 in the loop 50, and  $E_r$  represents a number in the range of 0 to  $(N-1)$ ). In the current embodiment, each PE is assigned a different  $E_r$  value for controlling the rounding. The simplest form for the function  $E$  is the case in which  $E_r = P_r$ , where  $P_r$  represents the PEs position in the loop. For example, for  $PE_0$ ,  $E_0 = 0$ ; for  $PE_1$ ,  $E_1 = 1$ ; for  $PE_2$ ,  $E_2 = 2$ ; etc. By assigning each PE in the loop a different  $E_r$  value, the rounding function can be controlled such that some of the local means are rounded up and some of the local means are rounded

down, thus insuring that  $V = \sum_{i=0}^{N-1} M_i$ . It should be noted that in the current embodiment, the local mean for each PE 30 in the loop is computed in parallel with the local means of the other PEs in the loop.

[0045] Table 1 illustrates the local mean calculation for the loop 50 as illustrated in FIG. 3 in which the total number of tasks on the loop is equal to forty-three (43). Referring to Table 1, it is apparent that the rounding function controls the rounding such that  $M_0$  through  $M_4$  are all rounded to five (5), whereas  $M_5$  through  $M_7$  are all rounded to six (6). The sum of the values of  $M_0$  through  $M_7$  is equal to forty-three (43), which equals the total number of tasks ( $V$ ) on the loop. Thus, tasks are neither lost nor gained due to rounding.

$PE_r$	$v_r$	$E_r$	$(V+E_r)/N$	$M_r = Trunc((V+E_r)/N)$	$D_r$
$PE_0$	3	0	5.375	5	-2
$PE_1$	6	1	5.5	5	1
$PE_2$	2	2	5.625	5	-3
$PE_3$	7	3	5.75	5	2
$PE_4$	8	4	5.875	5	3
$PE_5$	5	5	6	6	-1
$PE_6$	5	6	6.125	6	-1
$PE_7$	7	7	6.25	6	1

Table #1 – Local Mean Calculation for the Loop 50 ( $V = 43$ ,  $N = 8$ ).

[0046] After the local means are computed in operation 62, the local deviation  $D_r$  is calculated for each PE in operation 63. In the current embodiment, the local deviation is simply the difference between the local number of tasks and the local mean (i.e.,  $D_r = v_r - M_r$ ). The local deviations for  $PE_0$  through  $PE_7$  are illustrated in Table #1.

[0047] After the local deviations are computed in operation 63, the sum weighted deviation in the anti-clockwise half of loop ( $A$ ) is determined for each PE in operation 64. The anti-clockwise sum ( $A$ ) is then formed in a similar manner as that used to form the partial value sum ( $V$ ) in operation 61. In operation 64, however, a weighing factor ( $w_r$ ) is assigned to each PE and the local weighted deviations ( $w_r D_r$ ) are then rotated halfway around the loop in clockwise direction and summed. In the current embodiment, greater weight is given to those PEs that are located closer to the selected PE (i.e., PEs that are closer to the selected PE have a greater weighing factor ( $w_r$ )). For example if  $PE_2$  is the selected element, then weighing factors are assigned to  $PE_1$ ,  $PE_0$ , and  $PE_7$  such that  $w_1 > w_0 > w_7$ . The sum weighted

deviation in the anti-clockwise half of loop can be represented by the equation:

$$A = \sum_{i=1}^{i=(N/2)-1} w_i D_i .$$

[0048] After the sum weighted deviation in the anti-clockwise half of loop ( $A$ ) is determined in operation 64, the sum weighted deviation in the clockwise half of loop ( $C$ ) are determined for each PE in operation 65. The clockwise sum ( $C$ ) is formed in a similar manner as that used to determine the anti-clockwise sum ( $A$ ) in operation 64. In operation 65, however, the local weighted deviations ( $w_i D_i$ ) are rotated halfway around the loop in an anti-clockwise direction and summed. As discussed in conjunction with operation 64, greater weight is given to those PEs that are located closer to the selected PE (i.e., PEs that are closer to the selected PE have a greater weighing factor ( $w_i$ )). Again if  $PE_2$  is the selected element, then weighing factors are assigned to  $PE_3$ ,  $PE_4$ , and  $PE_5$  such that  $w_3 > w_4 > w_5$ . The sum deviation in the clockwise half of loop can be represented by the equation:

$$C = \sum_{i=(N/2)+1}^{i=N-1} w_{-i} D_{-i} .$$

[0049] FIG. 5 illustrates how the sum weighted deviation in the anti-clockwise half of loop ( $A$ ) and the sum weighted deviation in the clockwise half of loop ( $C$ ) is determined for  $PE_2$ . As seen in FIG. 5, the sum weighted deviation in the clockwise half of loop ( $C$ ) is determined by combining  $PE_3$ ,  $PE_4$ , and  $PE_5$  into a "super PE". The sum weighted deviation of this super PE is  $C = \text{Sum}(w_3 D_3 + w_4 D_4 + w_5 D_5)$ . Likewise, the sum weighted deviation in the anti-clockwise half of loop ( $A$ ) is determined by combining  $PE_1$ ,  $PE_0$ , and  $PE_7$  into another "super PE". The sum deviation of this super PE is  $A = \text{Sum}(w_1 D_1 + w_0 D_0 + w_7 D_7)$ . It should be noted that in the current embodiment no weight is given to  $PE_6$ .

[0050] Referring to Table #1, the sum deviation in the clockwise half of loop ( $C$ ) using this super PE is  $w_3 D_3 + w_4 D_4 + w_5 D_5 = w_3(2) + w_4(3) + w_5(-1)$ . If weighing factors are assigned to  $PE_3$ ,  $PE_4$ , and  $PE_5$  as discussed above, for example in the current embodiment,  $w_3 = 3$ ,  $w_4 = 2$ , and  $w_5 = 1$ , then  $C = 3(2) + 2(3) + 1(-1) = 11$ . Likewise, the sum deviation in the anti-clockwise half of loop ( $A$ ) using the other super PE is  $w_1 D_1 + w_0 D_0 + w_7 D_7 = w_1(1) + w_0(-2) + w_7(1)$ . Again, if weighing factors are assigned to  $PE_1$ ,  $PE_0$ , and  $PE_7$  as discussed above, for example in the current embodiment,  $w_1 = 3$ ,  $w_0 = 2$ , and  $w_7 = 1$ , then  $A = 3(1) + 2(-2) + 1(1) = 0$ .

[0051] After the sum deviation in the clockwise half of loop ( $C$ ) is determined in operation 64 and the sum deviation in the anti-clockwise half of loop ( $A$ ) is determined in operation 65, clockwise and anti-clockwise transfer parameters ( $T_c$  and  $T_a$ , respectively) are determined in operation 66. Referring again to FIG. 5 from the perspective of  $PE_2$ , the loop has four values  $C$ ,  $A$ ,  $O$ , and  $S$ , where  $C$  represents both the sum deviation in the clockwise half of loop and

the deviation of the first “super PE”,  $A$  represents both the sum deviation in the anti-clockwise half of loop and the deviation of the second “super PE”,  $S$  represents the deviation of the selected PE (e.g., here  $PE_2$ ) and  $O$  represents the deviation of the PE opposite to the selected PE (i.e., the PE for which the local deviation is being determined; here,  $PE_2$ ). The selected PE can deduce the deviation value of its opposite PE) because all deviations in the loop must sum to zero (i.e.,  $A + C + S + O = 0$ ). It should be noted that  $A$  and  $C$  are calculated for each PE in parallel.

[0052] It should be noted that in the current embodiment, the weights assigned to each PE are selected such that a linear relationship exist between the weights and each PE’s location around the loop. Thus for example, through the use of an intermediate sum ( $K$ ), the weighted sums in the anti-clockwise direction ( $A$ ) can be calculated without using multiplication. Initially,  $K_0 = A_0 = D_0$  and as each value  $D_i$  ( $i = 1$  to  $N - 1$ ) is rotated through the local PE the a calculation for  $K_i$  (e.g.,  $K_i = K_{i-1} + D_i$ ) and a calculation for  $A_i$  (e.g.,  $A_i = A_{i-1} + K_i$ ) is performed. After ‘ $r$ ’ deviations have been rotated, the values of  $K_r$  and  $A_r$  are given by the following equations:

$$K_r = \sum_{i=0}^{i=r} D_i \text{ and;}$$

$$A_r = \sum_{i=0}^{i=r} K_i$$

[0053] It should be noted that the same strategy can be used for evaluating the weighted clockwise sum  $C$ . Using the above system of weighting,  $T_c$  is determined from the equation  $T_c = (S/4) + \Delta$  and  $T_a$  is determined from the equation  $T_a = (S/4) - \Delta$ , where  $\Delta = (A - C)/4N$ . In the current embodiment,  $\Delta = (A - C)/4N = (0 - 11)/32 = 0.34375$  and thus,  $T_c = (-3/4) + 0.34375 = 0.40625$ , and  $T_a = (-3/4) - 0.34375 = -1.09375$ .

[0054] It should be noted that the values obtained for  $T_c$  and  $T_a$  may need to be rounded in such a manner that  $R(T_c) + R(T_a) = D_r$ . In the current embodiment, tasks are transmitted in only one direction at a time around the loop (i.e., either in the clockwise or anti-clockwise direction). A direction is selected for the ‘first’ transmission around the loop and the values for  $T_c$  and  $T_a$  are rounded up in this direction. It should be noted that by ensuring ‘excess traffic’ is sent in the ‘first’ direction, the chance of the process finishing one step earlier is increased. In the current embodiment, tasks are transmitted in the anti-clockwise first, such that  $R(T_a) = \text{Ceil}(T_a)$ , where the ‘Ceil’ function returns the closest integer greater than or equal to the supplied input. To ensure that extra tasks are not created or lost by the rounding of  $R(T_a)$ ,  $R(T_c)$  is set equal  $D_r - R(T_a)$ .

[0055] Accordingly in the example above, tasks are transmitted anti-clockwise first such that  $T_a = -1.09375$  is rounded up to -1. To ensure that extra tasks are not created or lost by the

rounding of  $R(T_a)$ ,  $R(T_c)$  is set equal  $D_r - R(T_a)$ . Thus,  $T_c$  is equal to -2 [i.e.,  $-3 - (-1) = -2$ ]. It should be noted that other rounding mechanisms may be used while remaining within the scope of the present invention. For example,  $T_c$  may be rounded up on odd numbered PE's and  $T_a$  rounded up on even numbered PE's such that pairs of odd and even PE's exchange their 'excess traffic'.

[0056] In the case where the loop 50 is comprised of an odd number of PEs 30, an extra "phantom" PE may be used. The phantom PE is assigned a deviation of zero and is located diametrically opposite from the perspective of the selected PE (i.e., the PE for which the local deviation is being determined). For example, assume that loop 50 only has seven PEs (i.e.,  $PE_0$  to  $PE_6$ ). To calculate the local deviation of  $PE_0$ , the phantom PE would be placed between  $PE_3$  and  $PE_4$ ; for  $PE_1$ , between  $PE_4$  and  $PE_5$ ; for  $PE_2$ , between  $PE_5$  and  $PE_6$ , etc. Thus, the number of PEs between the selected PE and the phantom PE in the clockwise direction is equal to the number of PEs between the selected PE and the phantom PE in the anti-clockwise direction.

[0057] In some instances, it may be desirable to clamp the transfer rates to reduce the number of iterations needed to balance the loop. In one embodiment, a non-linear clamping operation is utilized. For example, the equations for  $T_c$  and  $T_a$  may be re-written as  $T_c = Trunc[(2S + \Delta) \div 4]$  and  $T_a = Trunc[(2S - \Delta) \div 4]$ , respectively, where  $\Delta = (A - C)$  represents the number of 'thru' tasks (i.e., the number of tasks passing through the current PE. If  $T_c$  and  $T_a$  are of opposite sign, then the number of 'thru' values may be reduced by clamping either  $T_c$  or  $T_a$  to zero. The remaining value (i.e.,  $T_a$  or  $T_c$  respectively) may then be found using the identity  $S = T_a + T_c$ . This ensures that any rounding error introduced by the *Trunc* function is correctly compensated for such that  $S$  is finally equal to zero.

[0058] In the current embodiment for example, the transfer parameter with the smallest absolute magnitude may be selected. The desired result can be achieved by applying the following non-linear modifications to  $\Delta$  where  $Mag = abs(2S)$ . If  $\Delta > Mag$  then set  $\Delta$  equal to  $Mag$  and if  $\Delta < -Mag$ , then set  $\Delta$  equal to  $-Mag$ . The revised value for  $\Delta$  (i.e.,  $\Delta = Mag$  or  $\Delta = -Mag$ ) is then substituted into the equations  $T_c = Trunc[(2S + \Delta) \div 4]$  and  $T_a = Trunc[(2S - \Delta) \div 4]$ . It should be noted that other clamping operations may be used while remaining within the scope of the present invention.

[0059] After the clockwise and anti-clockwise transfer parameters are determined in operation 66, the tasks are redistributed among the PEs in response to the clockwise and anti-clockwise transfer parameters (i.e.,  $T_c$  and  $T_a$ , respectively) in operation 67. In the current embodiment, a positive  $T_c$  parameter represents the number of values that are to be transmitted clockwise out of the local PE. A negative  $T_c$  parameter represents the number of values that are to be transmitted from the clockwise PE into the local PE. Similarly, a

positive  $T_a$  parameter represents the number of values that are to be transmitted anti-clockwise out of the local PE. A negative  $T_a$  parameter represents the number of values that are to be transmitted from the anti-clockwise PE into the local PE.

[0060] If the local deviation ( $D$ ) is negative, one or more of the received values will be "absorbed" by the local PE to make up the local deficit. The other will be transmitted, either from the clockwise PE to the anti-clockwise PE, or from the anti-clockwise PE to the clockwise PE. On occasion, some PEs may start off with no values at all, these PEs may have to "mark time" until they receive a value. It should be noted that after each successful transmission or receipt, the local parameters  $T_c$  and  $T_a$  need to be updated. The redistribution stage only terminates when  $T_c = T_a = 0$  for all PEs.

[0061] As discussed above, the clockwise transfer parameter for  $PE_2$  in the current embodiment is  $T_c = -2$ . Because  $T_c$  is negative, 2 tasks are to be transmitted from  $PE_3$  into  $PE_2$ . Likewise, the anti-clockwise transfer parameter for  $PE_2$  in the current embodiment is  $T_a = -1$ . Because  $T_a$  is negative, one task is to be transmitted from  $PE_1$  into  $PE_2$ . It should be apparent that  $PE_2$  had a deviation of  $-3$  (i.e.,  $D_2 = -3$ ). Thus, three tasks were transferred into  $PE_2$  in operation 67.

[0062] It should be recognized that the above-described embodiments of the invention are intended to be illustrative only. Numerous alternative embodiments may be devised by those skilled in the art without departing from the scope of the following claims.

What is claimed is:

1. A method for balancing the load of a parallel processing system having a plurality of parallel processing elements arranged in a loop, wherein each processing element has a local number of tasks associated therewith, wherein  $r$  represents the number for a selected processing element  $PE_r$ , and wherein each of said processing elements are operable to communicate with a clockwise adjacent processing element and with an anti-clockwise adjacent processing element, the method comprising:
  - determining within each of said processing elements a total number of tasks present within said loop;
  - calculating a local mean number of tasks within each of said plurality of processing elements;
  - calculating a local deviation within each of said plurality of processing elements;
  - determining a sum weighted deviation within each of said processing elements for one-half of said loop in an anti-clockwise direction, said one-half of said loop being relative to each of said selected processing elements;
  - determining a sum weighted deviation within each of said processing elements in one-half of said loop in a clockwise direction, said one-half of said loop being relative to each of said selected processing element;
  - determining a clockwise transfer parameter and an anti-clockwise transfer parameter within each of said processing elements; and
  - redistributing tasks among said plurality of processing elements in response to said clockwise transfer parameters and said anti-clockwise parameters within each of said plurality of processing elements.
2. The method of claim 1 wherein said determining within each of said processing elements a total number of tasks present within said loop, comprises:
  - transmitting said local number of tasks associated with each of said processing elements to each other of said plurality of processing elements within said loop;
  - receiving within each of said processing elements said number of local tasks associated with said each other of said plurality of processing elements; and
  - summing said number of local tasks associated with each of said processing elements with said number of local tasks associated with each other of said plurality of processing elements.

3. The method of claim 1 wherein said determining said total number of tasks present within said loop includes solving the equation  $V = \sum_{i=-N}^{i=N-1} v_i$ , where  $V$  represents said total number of tasks,  $2N$  represents the number of processing elements in said loop, and  $v_i$  represents said local number of tasks associated with an  $i^{th}$  processing element in said loop.
4. The method of claim 1 wherein said calculating a local mean number of tasks within each of said plurality of processing elements ( $PE_r$ ) includes solving the equation  $M_r = Trunc((V + E_r) / 2N)$ , where  $M_r$  is said local mean for  $PE_r$ , where  $2N$  is the total number of processing elements in said loop, and where  $E_r$  is a number in the range of 0 to  $(2N-1)$  and wherein each processing element has a different  $E_r$  value.
5. The method of claim 4 wherein  $E_r$  controls said *Trunc* function such that said total number of tasks ( $V$ ) for said loop is equal to the sum of the local mean number of tasks ( $M_r$ ) for each of said plurality of processing elements in said loop (i.e.,  $V = \sum_{i=-N}^{i=N-1} M_i$ ).
6. The method of claim 4 wherein said local mean  $M_r = Trunc((V + E_r) / N)$  for each local  $PE_r$  within said loop is equal to one of  $X$  and  $(X+1)$ .
7. The method of claim 1 wherein said calculating a local deviation within each of said plurality of processing elements comprises finding the difference between said local number of tasks and said local mean number for each of said plurality of processing elements.
8. The method of claim 1 wherein said determining a sum weighted deviation within each of said processing elements for one-half of said loop in an anti-clockwise direction comprises:
  - assigning a weight to each other of said plurality of processing elements within said loop;
  - transmitting said local deviation and said weight associated with each of said processing elements half way around said loop in an anti-clockwise direction, said one-half of said loop being relative to each of said selected processing elements;



receiving said local deviation and said weight associated with each other of said processing elements half way around said loop in a clockwise direction, said one-half of said loop being relative to each of said selected processing elements; and

summing the product of said local deviation and said weight associated with each other of said processing elements half way around said loop in a clockwise direction.

9. The method of claim 1 wherein said determining a sum weighted deviation within each of said processing elements in one-half of said loop in a clockwise direction comprises:

assigning a weight to each other of said plurality of processing elements within said loop;

transmitting said local deviation and said weight associated with each of said processing elements half way around said loop in an clockwise direction, said one-half of said loop being relative to each of said selected processing elements;

receiving said local deviation and said weight associated with each other of said processing elements half way around said loop in a anti-clockwise direction, said one-half of said loop being relative to each of said selected processing elements; and

summing the product of said local deviation and said weight associated with each other of said processing elements half way around said loop in a anti-clockwise direction.

10. The method of claim 1 wherein said determining a clockwise transfer parameter and an anti-clockwise transfer parameter within each of said processing elements comprises:

setting  $T_a = (S/4) - \Delta$ ; and

setting  $T_c = (S/4) + \Delta$ , where  $T_c$  represents said clockwise transfer parameter,  $T_a$  represents said anti-clockwise transfer parameter,  $\Delta = (A - C)/4N$ ,  $A$  represents the sum weighted deviation within each of said processing elements in one-half of said loop in an anti-clockwise direction,  $C$  represents sum weighted deviation within each of said processing elements in one-half of said loop in a clockwise direction, and  $N$  represents the number of PEs on the loop.

11. The method of claim 1 wherein said determining a clockwise transfer parameter and an anti-clockwise transfer parameter within each of said processing elements comprises at least one of:

setting  $T_c = \text{Trunc}[(2S + \Delta) \div 4]$  and  $T_a = S - T_c$ ; and

setting the  $T_a = \text{Trunc}[(2S - \Delta) \div 4]$  and  $T_c = S - T_a$ ;

where  $T_c$  represents said clockwise transfer parameter, where  $T_a$  represents said anti-clockwise transfer parameter, where  $\Delta = \text{Mag}$ , if  $\Delta > \text{Mag}$ , where  $\Delta = -\text{Mag}$ , if  $\Delta < -\text{Mag}$ , where  $\text{Mag} = \text{abs}(2S)$ , and where  $S$  represents the local deviation of a selected processing element.

12. A method for reassigning tasks among an odd numbered plurality of processing elements within a parallel processing system, said processing elements being connected in a loop and each having a local number of tasks associated therewith, the method comprising:

- determining a total number of tasks on said loop;
- computing a local mean value for a selected processing element;
- computing a local deviation for said selected processing element, said local deviation representative of the difference between said local number of tasks for said selected processing element and said local mean value for said selected processing element;
- inserting a phantom processing element within said loop;
- assigning a weight to each of said plurality of processing elements;
- summing a weighted deviation of said processing elements located within one-half of the loop in an anti-clockwise direction relative to said selected processing element;
- summing said weighted deviation of said processing elements located within one-half of the loop in a clockwise direction relative to said selected processing element;
- computing a number of tasks to transfer in a clockwise direction for said selected processing element;
- computing a number of tasks to transfer in an anti-clockwise direction for said selected processing element; and
- reassigning tasks relative to the said number of tasks to transfer in a clockwise direction and said number of task to transfer in an anti-clockwise direction.

13. The method of claim 12 wherein said determining the total number of tasks on said loop, comprises:

- transmitting said local number of tasks associated with each of said processing elements to each other of said plurality of processing elements within said loop;
- receiving within each of said processing elements said number of local tasks associated with said each other of said plurality of processing elements; and

summing said number of local tasks associated with each of said processing elements with said number of local tasks associated with each other of said plurality of processing elements.

14. The method of claim 12 wherein computing a local mean value for a selected processing element includes solving the equation  $M_r = \text{Trunc}((V + E_r) / 2N)$ , where  $M_r$  is said local mean for a selected PE<sub>r</sub>,  $2N$  is the total number of processing elements in said loop, and  $E_r$  is a number in the range of 0 to  $(2N - 1)$  and wherein each processing element has a different  $E_r$  value.

15. The method of claim 14 wherein  $E_r$  controls said *Trunc* function such that said total number of tasks ( $V$ ) for said loop is equal to the sum of the local mean number of tasks ( $M_r$ ) for each of said plurality of processing elements in said loop (i.e.,

$$V = \sum_{i=0}^{i=2N-1} M_i).$$

16. The method of claim 12 wherein said inserting a phantom processing element within said loop further comprises:

locating said phantom processing element in a position within said loop that is diametrically opposed to said processing element; and  
assigning a zero deviation value to said phantom processing element.

17. The method of claim 12 wherein said assigning a weight to each of said plurality of processing elements includes assigning a weight dependent upon each of said processing element's location to said selected processing element.

18. The method of claim 12 wherein said computing a local mean value for a selected processing element, said computing a local deviation for said selected processing element, said inserting a phantom processing element within said loop, said assigning a weight to each of said plurality of processing elements, said summing said weighted deviation of said processing elements located within one-half of the loop in an anti-clockwise direction, summing said weighted deviation of said processing elements located within one-half of the loop in a clockwise direction, computing a number of tasks to transfer in a clockwise direction for said selected processing element, computing a number of tasks to transfer in an anti-clockwise direction for said selected processing element, and reassigning tasks relative to the said number of tasks

to transfer in a clockwise direction and said number of tasks to transfer in an anti-clockwise direction are completed simultaneously for each of said plurality of processing elements within said loop.

19. The method of claim 12 wherein said summing said weighted deviation of said processing elements located within one-half of the loop in an anti-clockwise direction relative to said selected processing element comprises:

transmitting said local weighted deviation associated with each of said processing elements half way around said loop in an anti-clockwise direction, said one-half of said loop being relative to each of said selected processing elements;

receiving said local weighted deviation associated with each other of said processing elements half way around said loop in a clockwise direction, said one-half of said loop being relative to each of said selected processing elements; and

summing said local weighted deviations associated with each other of said processing elements half way around said loop in a clockwise direction.

20. The method of claim 12 wherein summing said weighted deviation of said processing elements located within one-half of the loop in a clockwise direction relative to said selected processing element comprises:

transmitting said local weighted deviation associated with each of said processing elements half way around said loop in an clockwise direction, said one-half of said loop being relative to each of said selected processing elements;

receiving said local weighted deviation associated with each other of said processing elements half way around said loop in a anti-clockwise direction, said one-half of said loop being relative to each of said selected processing elements; and

summing said local weighted deviations associated with each other of said processing elements half way around said loop in an anti-clockwise direction.

21. A memory device carrying a set of instructions which, when executed, perform a method comprising:

determining within each of said processing elements a total number of tasks present within said loop;

calculating a local mean number of tasks within each of said plurality of processing elements;

calculating a local deviation within each of said plurality of processing elements;

determining a sum weighted deviation within each of said processing elements for one-half of said loop in an anti-clockwise direction, said one-half of said loop being relative to each of said selected processing elements;

determining a sum weighted deviation within each of said processing elements in one-half of said loop in a clockwise direction, said one-half of said loop being relative to each of said selected processing element;

determining a clockwise transfer parameter and an anti-clockwise transfer parameter within each of said processing elements; and

redistributing tasks among said plurality of processing elements in response to said clockwise transfer parameters and said anti-clockwise parameters within each of said plurality of processing elements.

## ABSTRACT OF THE DISCLOSURE

One aspect of the present invention relates to a method for balancing the load of a parallel processing system having a plurality of parallel processing elements arranged in a loop, wherein each processing element has a local number of tasks associated therewith. The method comprises determining within each processing element a total number of tasks present within the loop, calculating a local mean number of tasks within each processing element, assigning a weight to each of said plurality of processing elements, and calculating a local weighted deviation within each processing element. The method also comprises determining the sum weighted deviations within each processing element for one-half the loop in an anti-clockwise direction and in a clockwise direction, determining clockwise and anti-clockwise transfer parameters within each processing element, and redistributing tasks among the processing elements in response to the clockwise and anti-clockwise transfer parameters.

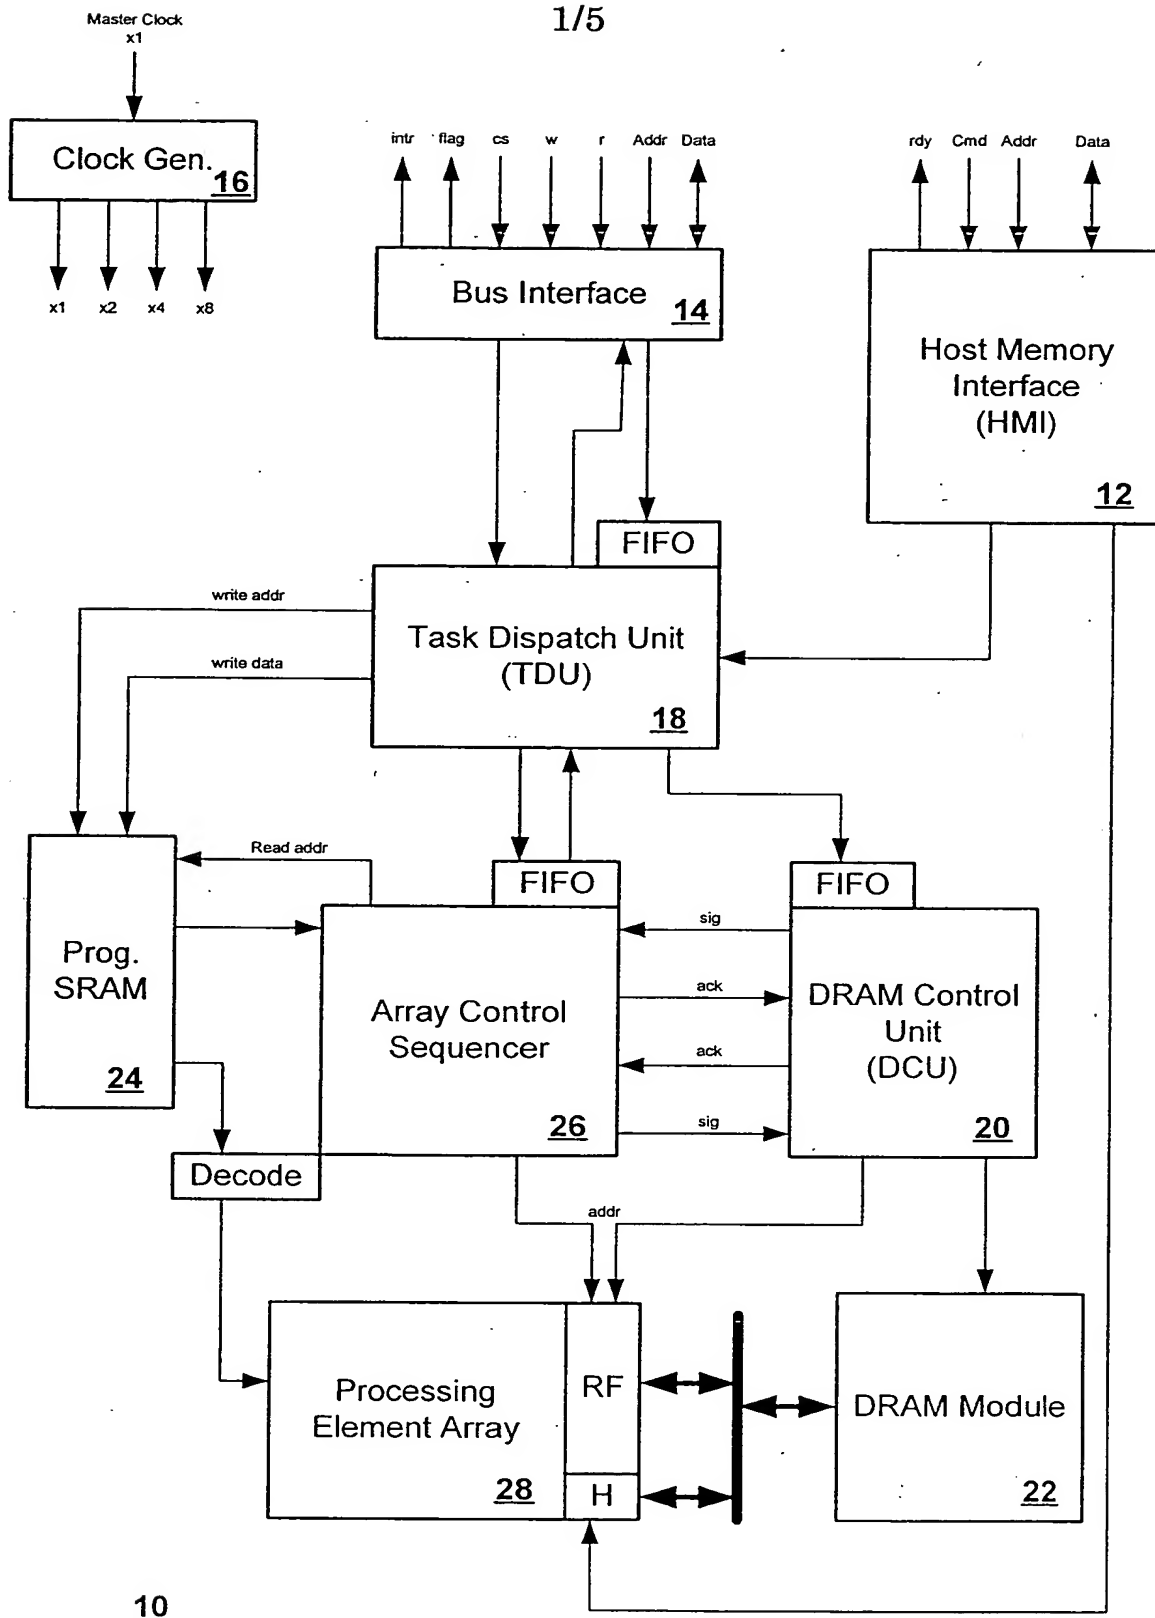


FIG 1





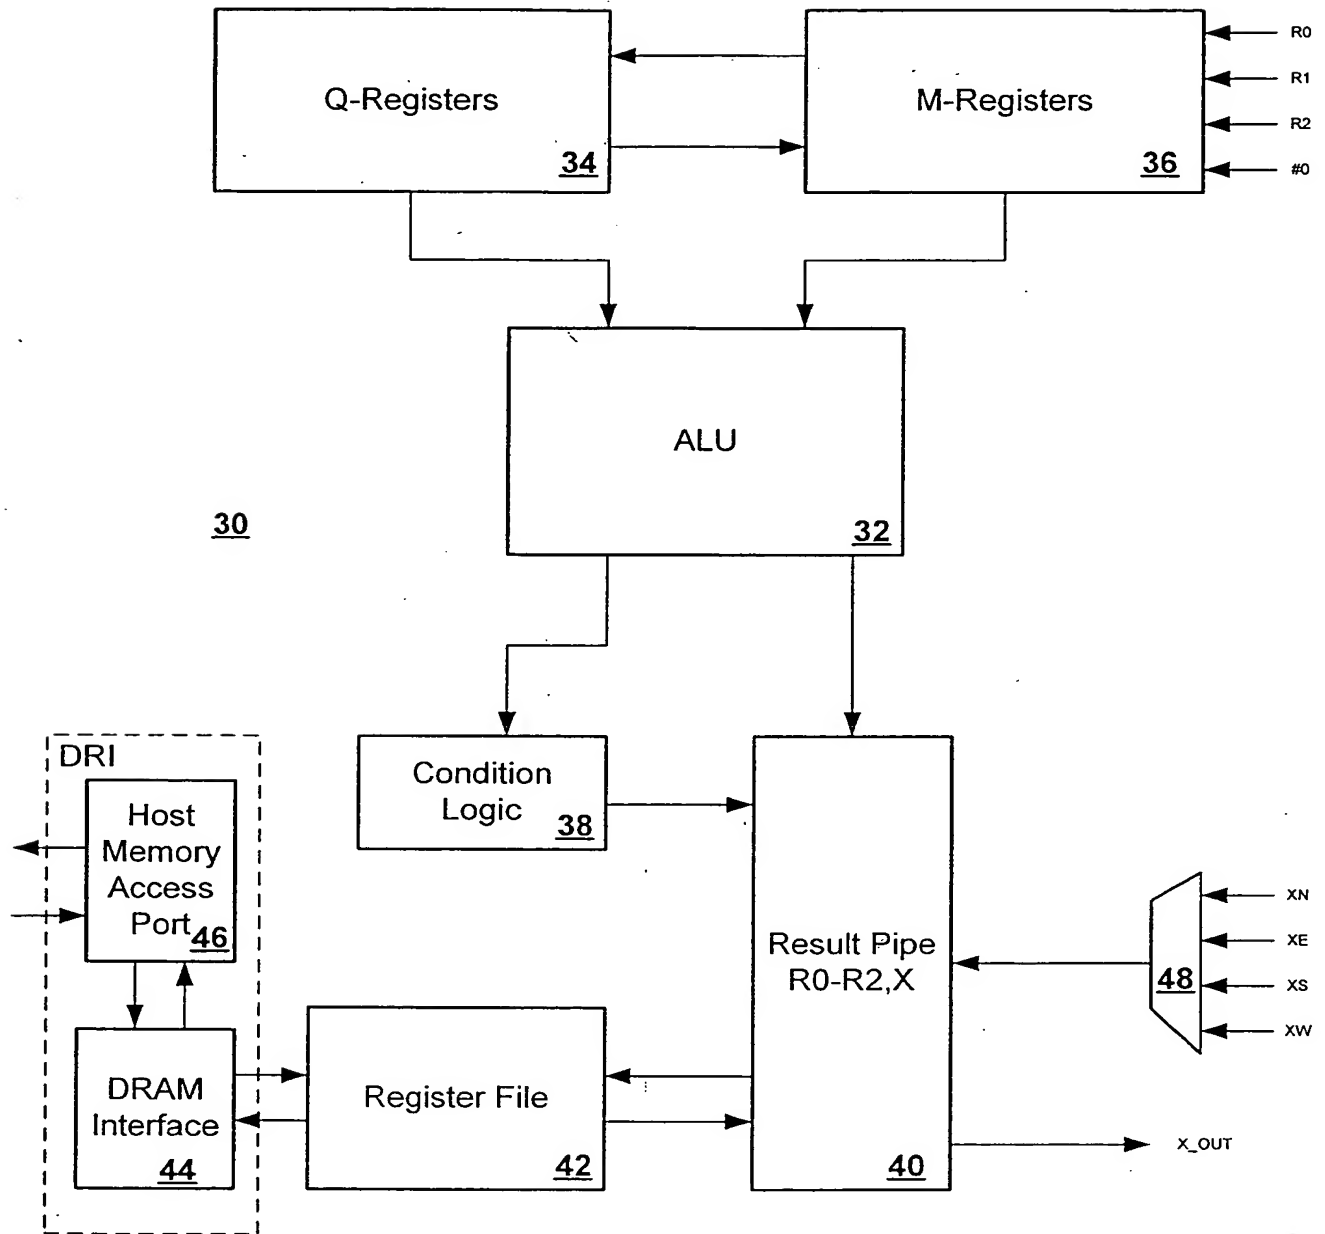


FIG 2






FIG 3



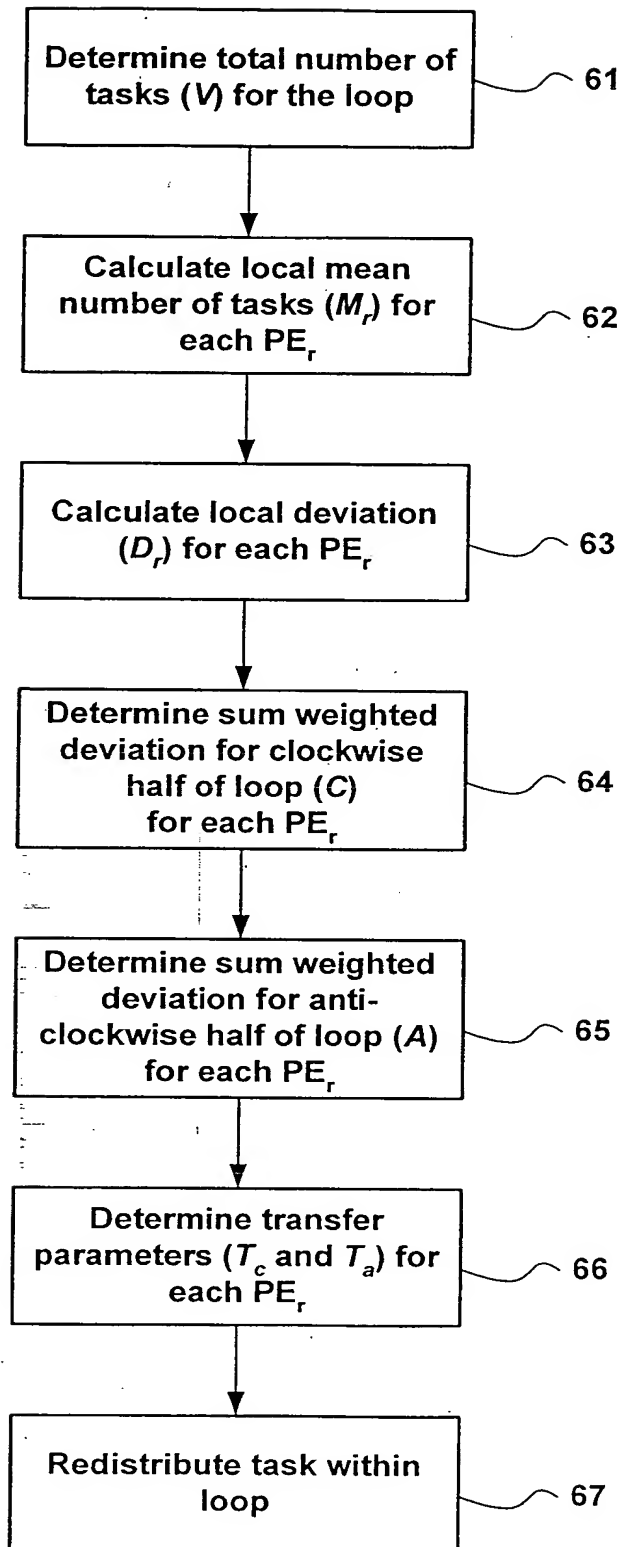


FIG 4



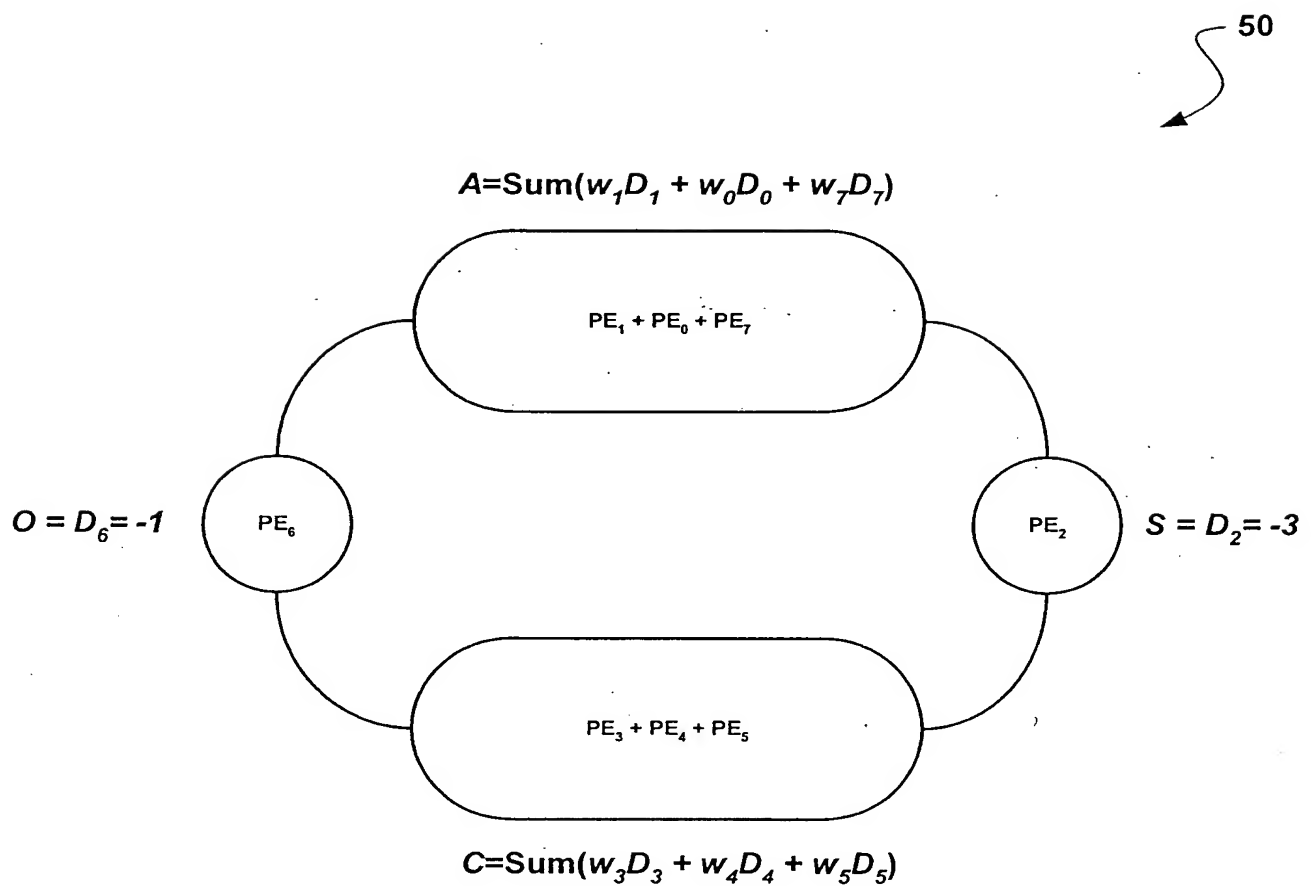


FIG 5

